Virtual GPU Software R510 for Microsoft Windows Server

Release Notes

RN-08686-001 _v14.0 through 14.4 | December 2022

Table of Contents

Chapter 1. Release Notes

These *Release Notes* summarize current status, information on validated platforms, and known issues with NVIDIA vGPU software and associated hardware on Microsoft Windows Server.

Note: The most current version of the documentation for this release of NVIDIA vGPU software can be found online at NVIDIA Virtual GPU Software Documentation.

1.1. NVIDIA vGPU Software Driver Versions

Each release in this release family of NVIDIA vGPU software includes a specific version of the NVIDIA Windows driver and NVIDIA Linux driver.

NVIDIA vGPU Software Version	NVIDIA Windows Driver Version	NVIDIA Linux Driver Version
14.4	514.08	510.108.03
14.3	513.91	510.108.03
14.2	513.46	510.85.02
14.1	512.78	510.73.08
14.0	511.65	510.47.03

Note:

All releases in this release family of NVIDIA vGPU software are compatible with **all** releases of the NVIDIA vGPU software license server.

1.2. Updates in Release 14.4

New Features in Release 14.4

- Security updates - see *Security Bulletin: NVIDIA GPU Display Driver - November 2022*, which is updated shortly after the release date of this software and is listed on the NVIDIA Product Security page

1.3. Updates in Release 14.3

New Features in Release 14.3

- Security updates - see *Security Bulletin: NVIDIA GPU Display Driver - November 2022*, which is posted shortly after the release date of this software and is listed on the NVIDIA Product Security page
- Support for non-transparent local proxy servers when NVIDIA vGPU software is served licenses by a Cloud License Service (CLS) instance
- Miscellaneous bug fixes

1.4. Updates in Release 14.2

New Features in Release 14.2

- Security updates - see *Security Bulletin: NVIDIA GPU Display Driver - August 2022*, which is posted shortly after the release date of this software and is listed on the NVIDIA Product Security page
- Miscellaneous bug fixes

Feature Support Withdrawn in Release 14.2

- All versions of Microsoft Windows Server 2016 with Hyper-V role are no longer supported as a hypervisor.
- All versions of Microsoft Windows Server 2016 are no longer supported as a guest OS.

1.5. Updates in Release 14.1

New Features in Release 14.1

- Security updates - see *Security Bulletin: NVIDIA GPU Display Driver - May 2022*, which is posted shortly after the release date of this software and is listed on the NVIDIA Product Security page
- Miscellaneous bug fixes

1.6. Updates in Release 14.0

New Features in Release 14.0

- Support for GPU System Processor (GSP) in GPU pass through and bare-metal configurations on Linux with vCS

  Note: If you are using a product other than vCS, you must disable GSP as explained in *Virtual GPU Software User Guide*.

- Addition of RPM and Debian packages for the NVIDIA vGPU software graphics drivers for Linux
- Security updates - see *Security Bulletin: NVIDIA GPU Display Driver - February 2022*, which is posted shortly after the release date of this software and is listed on the NVIDIA Product Security page
- Miscellaneous bug fixes

Hardware and Software Support Introduced in Release 14.0

- Support for the following GPUs:
  - NVIDIA A2
  - NVIDIA A30X
  - NVIDIA A100X
  - NVIDIA RTX A5500
- Support for Red Hat Enterprise Linux 8.5 as a guest OS
- Support for Red Hat Enterprise Linux 8.4 and 8.2 as a guest OS
- Support for Ubuntu 20.04 LTS as a guest OS

Feature Support Withdrawn in Release 14.0

- Red Hat Enterprise Linux 7.8 and 7.7 are no longer supported as a guest OS.

- Windows Server 2012 R2 is no longer supported as a guest OS.

Features Deprecated in Release 14.0

The following table lists features that are deprecated in this release of NVIDIA vGPU software. Although the features remain available in this release, they might be withdrawn in a future release. In preparation for the possible removal of these features, use the preferred alternative listed in the table.

Deprecated Feature	Preferred Alternative	Additional Information
Legacy NVIDIA vGPU software license server	NVIDIA License System	NVIDIA Virtual GPU Software License Server End of Life Notice

Chapter 2. Validated Platforms

This release family of NVIDIA vGPU software provides support for several NVIDIA GPUs on validated server hardware platforms, Microsoft Windows Server hypervisor software versions, and guest operating systems.

2.1. Supported NVIDIA GPUs and Validated Server Platforms

This release of NVIDIA vGPU software on Microsoft Windows Server provides support for several NVIDIA GPUs running on validated server hardware platforms. For a list of validated server platforms, refer to NVIDIA GRID Certified Servers.

The supported products for each type of NVIDIA vGPU software deployment depend on the GPU.

Note: All GPUs that support graphics acceleration are supported as a secondary device in a bare-metal deployment. Tesla M6 is also supported as the primary display device in a bare-metal deployment.

GPUs Based on the NVIDIA Ampere Architecture

GPU	Supported NVIDIA vGPU Software Products[1,2,3,4]		
	Time-Sliced NVIDIA vGPU	MIG-Backed NVIDIA vGPU	DDA
NVIDIA A100 PCIe 80GB	N/A	N/A	vCS
NVIDIA A100X	N/A	N/A	vCS
NVIDIA A100 HGX 80GB	N/A	N/A	vCS
NVIDIA A100 PCIe 40GB	N/A	N/A	vCS
NVIDIA A100 HGX 40GB	N/A	N/A	vCS
NVIDIA A40[5]	N/A	N/A	▶ vCS ▶ vWS

GPU	Supported NVIDIA vGPU Software Products[1,2,3,4]		
	Time-Sliced NVIDIA vGPU	MIG-Backed NVIDIA vGPU	DDA
			▶ vApps
NVIDIA A30	N/A	N/A	vCS
NVIDIA A30X	N/A	N/A	vCS
NVIDIA A16	N/A	N/A	▶ vCS ▶ vWS ▶ vApps
NVIDIA A10	N/A	N/A	▶ vCS ▶ vWS ▶ vApps
NVIDIA A2	N/A	N/A	▶ vCS ▶ vWS ▶ vApps
NVIDIA RTX A6000[5]	N/A	N/A	▶ vCS ▶ vWS ▶ vApps
NVIDIA RTX A5500[5]	N/A	N/A	▶ vCS ▶ vWS ▶ vApps
NVIDIA RTX A5000[5]	N/A	N/A	▶ vCS ▶ vWS ▶ vApps

GPUs Based on the NVIDIA Turing™ Architecture

GPU	Supported NVIDIA vGPU Software Products[1,2,3,4]		
	Time-Sliced NVIDIA vGPU	MIG-Backed NVIDIA vGPU	DDA
Tesla T4	N/A	N/A	▶ vCS ▶ vWS

GPU	Supported NVIDIA vGPU Software Products[1,2,3,4]		
	Time-Sliced NVIDIA vGPU	MIG-Backed NVIDIA vGPU	DDA
			▶ vApps
Quadro RTX 6000 5	N/A	N/A	▶ vCS ▶ vWS ▶ vApps
Quadro RTX 6000 passive5	N/A	N/A	▶ vCS ▶ vWS ▶ vApps
Quadro RTX 80005	N/A	N/A	▶ vCS ▶ vWS ▶ vApps
Quadro RTX 8000 passive5	N/A	N/A	▶ vCS ▶ vWS ▶ vApps

GPUs Based on the NVIDIA Volta Architecture

GPU	Supported NVIDIA vGPU Software Products[1,2,3,4]		
	Time-Sliced NVIDIA vGPU	MIG-Backed NVIDIA vGPU	DDA
Tesla V100 SXM2	N/A	N/A	▶ vCS ▶ vWS ▶ vApps
Tesla V100 SXM2 32GB	N/A	N/A	▶ vCS ▶ vWS ▶ vApps
Tesla V100 PCIe	N/A	N/A	▶ vCS ▶ vWS ▶ vApps
Tesla V100 PCIe 32GB	N/A	N/A	▶ vCS

GPU	Supported NVIDIA vGPU Software Products[1,2,3,4]		
	Time-Sliced NVIDIA vGPU	MIG-Backed NVIDIA vGPU	DDA
			▶ vWS ▶ vApps
Tesla V100S PCIe 32GB	N/A	N/A	▶ vCS ▶ vWS ▶ vApps
Tesla V100 FHHL	N/A	N/A	▶ vCS ▶ vWS ▶ vApps

GPUs Based on the NVIDIA Pascal™ Architecture

GPU	Supported NVIDIA vGPU Software Products[1,2,3,4]		
	Time-Sliced NVIDIA vGPU	MIG-Backed NVIDIA vGPU	DDA
Tesla P4	N/A	N/A	▶ vCS ▶ vWS ▶ vApps
Tesla P6	N/A	N/A	▶ vCS ▶ vWS ▶ vApps
Tesla P40	N/A	N/A	▶ vCS ▶ vWS ▶ vApps
Tesla P100 PCIe 16 GB	N/A	N/A	▶ vCS ▶ vWS ▶ vApps
Tesla P100 SXM2 16 GB	N/A	N/A	▶ vCS ▶ vWS ▶ vApps

GPU	Supported NVIDIA vGPU Software Products[1, 2, 3, 4]		
	Time-Sliced NVIDIA vGPU	MIG-Backed NVIDIA vGPU	DDA
Tesla P100 PCIe 12GB	N/A	N/A	▶ vCS ▶ vWS ▶ vApps

GPUs Based on the NVIDIA Maxwell™ Graphic Architecture

Note: NVIDIA Virtual Compute Server (vCS) is **not** supported on GPUs based on the NVIDIA Maxwell graphic architecture.

GPU	Supported NVIDIA vGPU Software Products[1, 2, 3, 4]		
	Time-Sliced NVIDIA vGPU	MIG-Backed NVIDIA vGPU	DDA
Tesla M6	N/A	N/A	▶ vWS ▶ vApps
Tesla M10	N/A	N/A	▶ vWS ▶ vApps
Tesla M60	N/A	N/A	▶ vWS ▶ vApps

[1] The supported products are as follows:

- vCS: NVIDIA Virtual Compute Server
- vWS: NVIDIA RTX Virtual Workstation
- vPC: NVIDIA Virtual PC
- vApps: NVIDIA Virtual Applications

[2] N/A indicates that the deployment is not supported.
[3] vCS is supported only on Linux operating systems.
[4] vApps is supported only on Windows operating systems.
[5] This GPU is supported only in displayless mode. In displayless mode, local physical display connectors are disabled.

2.1.1. Switching the Mode of a GPU that Supports Multiple Display Modes

Some GPUs support displayless and display-enabled modes but must be used in NVIDIA vGPU software deployments in displayless mode.

The GPUs listed in the following table support multiple display modes. As shown in the table, some GPUs are supplied from the factory in displayless mode, but other GPUs are supplied in a display-enabled mode.

GPU	Mode as Supplied from the Factory
NVIDIA A40	Displayless
NVIDIA RTX A5000	Display enabled
NVIDIA RTX A5500	Display enabled
NVIDIA RTX A6000	Display enabled

A GPU that is supplied from the factory in displayless mode, such as the NVIDIA A40 GPU, might be in a display-enabled mode if its mode has previously been changed.

To change the mode of a GPU that supports multiple display modes, use the `displaymodeselector` tool, which you can request from the NVIDIA Display Mode Selector Tool page on the NVIDIA Developer website.

Note:

Only the following GPUs support the `displaymodeselector` tool:

- NVIDIA A40
- NVIDIA RTX A5000
- NVIDIA RTX A5500
- NVIDIA RTX A6000

Other GPUs that support NVIDIA vGPU software do not support the `displaymodeselector` tool and, unless otherwise stated, do not require display mode switching.

2.1.2. Switching the Mode of a Tesla M60 or M6 GPU

Tesla M60 and M6 GPUs support compute mode and graphics mode. NVIDIA vGPU requires GPUs that support both modes to operate in graphics mode.

Recent Tesla M60 GPUs and M6 GPUs are supplied in graphics mode. However, your GPU might be in compute mode if it is an older Tesla M60 GPU or M6 GPU or if its mode has previously been changed.

To configure the mode of Tesla M60 and M6 GPUs, use the `gpumodeswitch` tool provided with NVIDIA vGPU software releases. If you are unsure which mode your GPU is in, use the `gpumodeswitch` tool to find out the mode.

Note:

Only Tesla M60 and M6 GPUs support the `gpumodeswitch` tool. Other GPUs that support NVIDIA vGPU do not support the `gpumodeswitch` tool and, except as stated in Switching the Mode of a GPU that Supports Multiple Display Modes, do not require mode switching.

Even in compute mode, Tesla M60 and M6 GPUs do **not** support NVIDIA Virtual Compute Server vGPU types.

For more information, refer to *gpumodeswitch User Guide*.

2.2. Hypervisor Software Releases

This release supports **only** the hypervisor software versions listed in the table.

Note: If a specific release, even an update release, is not listed, it's **not** supported.

Software	Version Supported	Notes
Microsoft Windows Server 2022	Windows Server 2022 with Hyper-V role	
Microsoft Windows Server 2019	Windows Server 2019 with Hyper-V role	
14.0, 14.1 only: Microsoft Windows Server 2016	Windows Server 2016 1803 with Hyper-V role Windows Server 2016 1709 with Hyper-V role Windows Server 2016 1607 with Hyper-V role	Not supported on the following GPUs: ▶ NVIDIA A100 HGX 80GB ▶ NVIDIA A100 PCIe 40GB ▶ NVIDIA A100 HGX 40GB

2.3. Guest OS Support

NVIDIA vGPU software supports several Windows releases and Linux distributions as a guest OS using GPU pass-through.

Microsoft Windows Server with Hyper-V role supports GPU pass-through over Microsoft Virtual PCI bus. This bus is supported through paravirtualized drivers.

Note:

Use only a guest OS release that is listed as supported by NVIDIA vGPU software with your virtualization software. To be listed as supported, a guest OS release must be supported not

only by NVIDIA vGPU software, but also by your virtualization software. NVIDIA **cannot** support guest OS releases that your virtualization software does not support.

NVIDIA vGPU software supports **only** 64-bit guest operating systems. No 32-bit guest operating systems are supported.

2.3.1. Windows Guest OS Support

NVIDIA vGPU software supports **only** the 64-bit Windows releases listed as a guest OS on Microsoft Windows Server.

Note:

If a specific release, even an update release, is not listed, it's **not** supported.

- Windows Server 2019
- **14.0, 14.1 only:** Windows Server 2016 1607, 1709
- Windows 11 21H2 (**not** supported on Microsoft Windows Server 2016)
- Windows 10 November 2021 Update (21H2) and all Windows 10 releases supported by Microsoft up to and including this release

2.3.2. Linux Guest OS Support

NVIDIA vGPU software supports **only** the 64-bit Linux distributions listed as a guest OS on Microsoft Windows Server.

Note:

If a specific release, even an update release, is not listed, it's **not** supported.

- **Since 14.3:** Red Hat Enterprise Linux 8.7
- **Since 14.1:** Red Hat Enterprise Linux 8.6
- **14.0 only:** Red Hat Enterprise Linux 8.5
- Red Hat Enterprise Linux 8.4
- **14.0 only:** Red Hat Enterprise Linux 8.2
- Red Hat Enterprise Linux 7.9
- **14.0 only:** CentOS Linux 8 (2111)
- **14.0 only:** CentOS Linux 8 (2004)
- CentOS Linux 8 (2105)
- CentOS 7.6-7.8
- Ubuntu 20.04 LTS
- Ubuntu 18.04 LTS
- Ubuntu 16.04 LTS

- SUSE Linux Enterprise Server 15 SP2
- SUSE Linux Enterprise Server 12 SP2

Chapter 3. Known Issues

3.1. NLS client fails to acquire a license with the error `The allowed time to process response has expired`

Description

A licensed client of NVIDIA License System (NLS) fails to acquire a license with the error `The allowed time to process response has expired`. This error can affect clients of a Cloud License Service (CLS) instance or a Delegated License Service (DLS) instance.

This error occurs when the time difference between the system clocks on the client and the server that hosts the CLS or DLS instance is greater than 10 minutes. A common cause of this error is the failure of either the client or the server to adjust its system clock when daylight savings time begins or ends. The failure to acquire a license is expected to prevent clock windback from causing licensing errors.

Workaround

Ensure that system clock time of the client and any server that hosts a DLS instance match the current time in the time zone where they are located.

To prevent this error from occurring when daylight savings time begins or ends, enable the option to automatically adjust the system clock for daylight savings time:

- **Windows:** Set the **Adjust for daylight saving time automatically** option.
- **Linux:** Use the `hwclock` command.

Status

Not a bug

Ref.

3859889

3.2. With multiple active sessions, **NVIDIA Control Panel** incorrectly shows that the system is unlicensed

Description

In an environment with multiple active desktop sessions, the **Manage License** page of **NVIDIA Control Panel** shows that a licensed system is unlicensed. However, the `nvidia-smi` command and the management interface of the NVIDIA vGPU software license server correctly show that the system is licensed. When an active session is disconnected and reconnected, the **NVIDIA Display Container** service crashes.

The **Manage License** page incorrectly shows that the system is unlicensed because of stale data in **NVIDIA Control Panel** in an environment with multiple sessions. The data is stale because **NVIDIA Control Panel** fails to get and update the settings for remote sessions when multiple sessions or no sessions are active in the VM. The **NVIDIA Display Container** service crashes when a session is reconnected because the session is not active at the moment of reconnection.

Status

Open

Ref.

3761243

3.3. VP9 and AV1 decoding with web browsers are not supported on Microsoft Windows Server 2019

Description

VP9 and AV1 decoding with web browsers are not supported on Microsoft Windows Server 2019 and later supported releases. This issue occurs because starting with Windows Server 2019, the required codecs are not included with the OS and are not available through the

Microsoft Store app. As a result, hardware decoding is not available for viewing YouTube videos or using collaboration tools such as Google Meet in a web browser.

Version

This issue affects Microsoft Windows Server releases starting with Windows Server 2019.

Status

Not an NVIDIA bug

Ref.

200756564

3.4. **NVIDIA Control Panel** is started only for the RDP user that logs on first

Description

On all supported Windows Server guest OS releases, **NVIDIA Control Panel** is started only for the RDP user that logs on first. Other users cannot start **NVIDIA Control Panel**. If more than one RDP user is logged on when **NVIDIA Control Panel** is started, it always opens in the session of the RDP user that logged on first, irrespective of which user started **NVIDIA Control Panel**. Furthermore, on Windows Server 2016, **NVIDIA Control Panel** crashes if a user session is disconnected and then reconnected while **NVIDIA Control Panel** is open.

Version

This issue affects all supported Windows Server guest OS releases.

Status

Open

Ref.

3334310

3.5. 14.0 Only: For some license deployments, GSP firmware remains enabled with unsupported products

Description

If GPU System Processor (GSP) firmware is enabled with an unsupported product in a GPU pass through or bare-metal configuration on Linux, the VM or bare-metal host should fail to acquire a license. This behavior is implemented to prevent the VM or host from being in an unsupported configuration. However, only license acquisition from a networked legacy NVIDIA vGPU software license server fails. For all other license deployments, the VM or host acquires a license and is in an unsupported configuration.

GSP is supported only for vCS in GPU pass through and bare-metal deployments on Linux. If you are using any other product in a GPU pass through or bare-metal deployment on Linux, you must disable the GSP firmware.

Note: For NVIDIA vGPU deployments on Linux and all NVIDIA vGPU software deployments on Windows, GSP is also not supported but GSP firmware is already disabled. For these deployments, this issue does not arise.

The GSP firmware might be enabled with an unsupported product in a GPU pass through or bare-metal configuration on Linux. In this situation, the following error message is written to the licensing event log file when the VM or host attempts to acquire a license:

```
Invalid feature requested for the underlying GSP firmware configuration.
Disable GSP firmware to use this feature.
```

For the location of the licensing event log file, refer to *Virtual GPU Client Licensing User Guide*.

Workaround

Ensure that the GSP firmware is disabled as explained in *Virtual GPU Software User Guide*.

Status

Resolved in NVIDIA vGPU software 14.1

Ref.

3523478

3.6. After an upgrade of the Linux graphics driver from an RPM package in a licensed VM, licensing fails

Description

After the NVIDIA vGPU software graphics driver for Linux is upgraded from an RPM package in a licensed VM, licensing fails. The `nvidia-smi vgpu -q` command shows the driver version and license status as N/A. Restarting the `nvidia-gridd` service fails with a `Unit not found` error.

Workaround

Perform a clean installation of the NVIDIA vGPU software graphics driver for Linux from an RPM package.

1. Remove the currently installed driver.
2. Install the new version of the driver.

```
$ rpm -iv nvidia-linux-grid-510_510.108.03_amd64.rpm
```

Status

Open

Ref.

3512766

3.7. After an upgrade of the Linux graphics driver from a Debian package, the driver is not loaded into the VM

Description

After the NVIDIA vGPU software graphics driver for Linux is upgraded from a Debian package, the driver is not loaded into the VM.

Workaround

Use one of the following workarounds to load the driver into the VM:

- Reboot the VM.
- Remove the `nvidia` module from the Linux kernel and reinsert it into the kernel.

  1. Remove the `nvidia` module from the Linux kernel.

     ```
     $ sudo rmmod nvidia
     ```

  2. Reinsert the `nvidia` module into the Linux kernel.

     ```
     $ sudo modprobe nvidia
     ```

Status

Not a bug

Ref.

200748806

3.8. A licensed client might fail to acquire a license if a proxy is set

Description

If a proxy is set with a system environment variable such as `HTTP_PROXY` or `HTTPS_PROXY`, a licensed client might fail to acquire a license.

Workaround

Perform this workaround on each affected licensed client.

1. Add the address of the NVIDIA vGPU software license server to the system environment variable `NO_PROXY`.

   The address must be specified exactly as it is specified in the client's license server settings either as a fully-qualified domain name or an IP address. If the `NO_PROXY` environment variable contains multiple entries, separate the entries with a comma (`,`).

   If high availability is configured for the license server, add the addresses of the primary license server and the secondary license server to the system environment variable `NO_PROXY`.

2. Restart the NVIDIA driver service that runs the core NVIDIA vGPU software logic.

   - On Windows, restart the **NVIDIA Display Container** service.
   - On Linux, restart the `nvidia-gridd` service.

Status

Closed

Ref.

200704733

3.9. Disconnected sessions cannot be reconnected or might be reconnected very slowly with NVWMI installed

Description

Disconnected sessions cannot be reconnected or might be reconnected very slowly when the NVIDIA Enterprise Management Toolkit (NVWMI) is installed. This issue affects Citrix Virtual Apps and Desktops and VMware Horizon sessions on Windows guest VMs.

Workaround

Uninstall NVWMI.

Status

Open

Ref.

3262923

3.10. **NVIDIA Control Panel** fails to start if launched too soon from a VM without licensing information

Description

If NVIDIA licensing information is not configured on the system, any attempt to start **NVIDIA Control Panel** by right-clicking on the desktop within 30 seconds of the VM being started fails.

Workaround

Restart the VM and wait at least 30 seconds before trying to launch **NVIDIA Control Panel**.

Status

Open

Ref.

200623179

3.11. Citrix Virtual Apps and Desktops session corruption occurs in the form of residual window borders

Description

When a window is dragged across the desktop in a Citrix Virtual Apps and Desktops session, corruption of the session in the form of residual window borders occurs.

Version

This issue affects only Citrix Virtual Apps and Desktops version 7 2003

Workaround

Use Citrix Virtual Apps and Desktops version 7 1912 or 2006.

Status

Not an NVIDIA bug

Ref.

200608675

3.12. On Linux, the frame rate might drop to 1 after several minutes

Description

On Linux, the frame rate might drop to 1 frame per second (FPS) after NVIDIA vGPU software has been running for several minutes. Only some applications are affected, for example, `glxgears`. Other applications, such as Unigine Heaven, are not affected. This behavior occurs because Display Power Management Signaling (DPMS) for the Xorg server is enabled by default and the display is detected to be inactive even when the application is running. When DPMS is enabled, it enables power saving behavior of the display after several minutes of inactivity by setting the frame rate to 1 FPS.

Workaround

1. If necessary, stop the Xorg server.

   ```
   # /etc/init.d/xorg stop
   ```

2. In a plain text editor, edit the `/etc/X11/xorg.conf file` to set the options to disable DPMS and disable the screen saver.

   a). In the `Monitor` section, set the `DPMS` option to `false`.

   ```
   Option "DPMS" "false"
   ```

   b). At the end of the file, add a `ServerFlags` section that contains option to disable the screen saver.

   ```
   Section "ServerFlags"
       Option "BlankTime" "0"
     EndSection
   ```

   c). Save your changes to `/etc/X11/xorg.conf file` and quit the editor.

3. Start the Xorg server.

   ```
   # etc/init.d/xorg start
   ```

Status

Open

Ref.

200605900

3.13. Microsoft DDA fails with some GPUs

Description

Microsoft Discrete Device Assignment (DDA) fails with GPUs that have more than 16 GB of GPU memory. After the NVIDIA vGPU software graphics driver is installed in the guest VM, a second display device appears on the GPU and the driver prompts for a reboot. After the reboot, the device disappears and the Microsoft Hyper-V Video device appears.

This issue occurs because less memory-mapped input/output (MMIO) space is configured for the operating system than the device requires.

Workaround

Perform this workaround in a **Windows Power Shell** window on the hypervisor host.

Set the upper MMIO space to the amount that the device requires to allow all of the MMIO to be mapped. Upper MMIO space starts at approximately 64 GB in address space.

```
Set-VM –HighMemoryMappedIoSpace mmio-space –VMName vm-name
```

mmio-space

The amount of MMIO space that the device requires, appended with the appropriate unit of measurement, for example, **64GB** for 64 GB of MMIO space.

The required amount of MMIO space depends on the amount of BAR1 memory on the installed GPUs and the number of GPUs assigned to the VM as follows:

mmio-space = 2 # *gpu-bar1-memory* # *assigned-gpus*

gpu-bar1-memory

The amount of BAR1 memory on one of the installed GPUs. For example, in a server in which eight GPUs are installed and each GPU has 32 GB of BAR1 memory, *gpu-bar1-memory* is 32 GB.

assigned-gpus

The number of GPUs assigned to the VM.

vm-name

The name of the VM to which the GPU is assigned.

The following example sets the upper MMIO space to 64 GB for the VM named `mygpuvm`, to which one GPU with 32 GB of BAR1 memory is assigned.

```
Set-VM –HighMemoryMappedIoSpace 64GB –VMName mygpuvm
```

For more information, see Deploy graphics devices using Discrete Device Assignment on the Microsoft technical documentation site.

Status

Not an NVIDIA bug

Ref.

2812853

3.14. DWM crashes randomly occur in Windows VMs

Description

Desktop Windows Manager (DWM) crashes randomly occur in Windows VMs, causing a blue-screen crash and the bug check `CRITICAL_PROCESS_DIED`. Computer Management shows problems with the primary display device.

Version

This issue affects Windows 10 1809, 1903 and 1909 VMs.

Status

Not an NVIDIA bug

Ref.

2730037

3.15. NVIDIA vGPU software graphics driver fails after Linux kernel upgrade with DKMS enabled

Description

After the Linux kernel is upgraded (for example by running `sudo apt full-upgrade`) with Dynamic Kernel Module Support (DKMS) enabled, the `nvidia-smi` command fails to run. If DKMS is enabled, an upgrade to the Linux kernel triggers a rebuild of the NVIDIA vGPU software graphics driver. The rebuild of the driver fails because the compiler version is incorrect. Any attempt to reinstall the driver fails because the kernel fails to build.

When the failure occurs, the following messages are displayed:

```
-> Installing DKMS kernel module:
       ERROR: Failed to run `/usr/sbin/dkms build -m nvidia -v  510.47.03 -k
 5.3.0-28-generic`:
       Kernel preparation unnecessary for this kernel. Skipping...
       Building module:
```

```
       cleaning build area...
       'make' -j8 NV_EXCLUDE_BUILD_MODULES='' KERNEL_UNAME=5.3.0-28-generic
IGNORE_CC_MISMATCH='' modules...(bad exit status: 2)
       ERROR (dkms apport): binary package for nvidia:  510.47.03 not found
       Error! Bad return status for module build on kernel: 5.3.0-28-generic
(x86_64)
       Consult /var/lib/dkms/nvidia/ 510.47.03/build/make.log for more information.
       -> error.
       ERROR: Failed to install the kernel module through DKMS. No kernel module
was installed;
       please try installing again without DKMS, or check the DKMS logs for more
information.
       ERROR: Installation has failed. Please see the file '/var/log/nvidia-
installer.log' for details.
       You may find suggestions on fixing installation problems in the README
available on the Linux driver download page at www.nvidia.com.
```

Workaround

When installing the NVIDIA vGPU software graphics driver with DKMS enabled, use one of the following workarounds:

- Before running the driver installer, install the `dkms` package, then run the driver installer with the `-dkms` option.
- Run the driver installer with the `--no-cc-version-check` option.

Status

Not a bug.

Ref.

2836271

3.16. Blue screen crash occurs or no devices are found after VM reset

Description

If a VM on Microsoft Windows Server with Hyper-V role is reset from the hypervisor host, a blue screen crash (BSOD) occurs on Windows VMs and the `nvidia-smi` command reports `No devices were found` on Linux VMs. This issue occurs only on Windows Server 2019 with Tesla T4 GPUs with SRIOV enabled, Quadro RTX 8000 passive GPUs, and Quadro RTX 6000 passive GPUs.

Workaround

Contact NVIDIA Enterprise Support for a workaround for this issue, referencing the knowledge base article *Workaround for Blue Screen Crashes On Hyper-V DDA With SRIOV-Enabled GPUs*. This article is available only to NVIDIA Enterprise Support personnel.

Status

Not an NVIDIA bug

Ref.

200567935

3.17. Frame capture while the interactive logon message is displayed returns blank screen

Description

Because of a known limitation with NvFBC, a frame capture while the interactive logon message is displayed returns a blank screen.

An NvFBC session can capture screen updates that occur after the session is created. Before the logon message appears, there is no screen update after the message is shown and, therefore, a black screen is returned instead. If the NvFBC session is created after this update has occurred, NvFBC cannot get a frame to capture.

Workaround

Press **Enter** or wait for the screen to update for NvFBC to capture the frame.

Status

Not a bug

Ref.

2115733

3.18. RDS sessions do not use the GPU with some Microsoft Windows Server releases

Description

When some releases of Windows Server are used as a guest OS, Remote Desktop Services (RDS) sessions do not use the GPU. With these releases, the RDS sessions by default use the Microsoft Basic Render Driver instead of the GPU. This default setting enables 2D DirectX applications such as Microsoft Office to use software rendering, which can be more efficient than using the GPU for rendering. However, as a result, 3D applications that use DirectX are prevented from using the GPU.

Version

- Windows Server 2019
- Windows Server 2016
- Windows Server 2012

Solution

Change the local computer policy to use the hardware graphics adapter for all RDS sessions.

1. Choose **Local Computer Policy** > **Computer Configuration** > **Administrative Templates** > **Windows Components** > **Remote Desktop Services** > **Remote Desktop Session Host** > **Remote Session Environment** .
2. Set the **Use the hardware default graphics adapter for all Remote Desktop Services sessions** option.

3.19. A segmentation fault in DBus code causes `nvidia-gridd` to exit on Red Hat Enterprise Linux and CentOS

Description

On Red Hat Enterprise Linux 6.8 and 6.9, and CentOS 6.8 and 6.9, a segmentation fault in DBus code causes the `nvidia-gridd` service to exit.

The `nvidia-gridd` service uses DBus for communication with **NVIDIA X Server Settings** to display licensing information through the **Manage License** page. Disabling the GUI for licensing resolves this issue.

To prevent this issue, the GUI for licensing is disabled by default. You might encounter this issue if you have enabled the GUI for licensing and are using Red Hat Enterprise Linux 6.8 or 6.9, or CentOS 6.8 and 6.9.

Version

Red Hat Enterprise Linux 6.8 and 6.9

CentOS 6.8 and 6.9

Status

Open

Ref.

- 200358191
- 200319854
- 1895945

3.20. No **Manage License** option available in **NVIDIA X Server Settings** by default

Description

By default, the **Manage License** option is not available in **NVIDIA X Server Settings**. This option is missing because the GUI for licensing on Linux is disabled by default to work around the issue that is described in A segmentation fault in DBus code causes nvidia-gridd to exit on Red Hat Enterprise Linux and CentOS.

Workaround

This workaround requires `sudo` privileges.

Note: Do **not** use this workaround with Red Hat Enterprise Linux 6.8 and 6.9 or CentOS 6.8 and 6.9. To prevent a segmentation fault in DBus code from causing the `nvidia-gridd` service from exiting, the GUI for licensing must be disabled with these OS versions.

If you are licensing a physical GPU for vCS, you **must** use the configuration file `/etc/nvidia/gridd.conf`.

1. If **NVIDIA X Server Settings** is running, shut it down.
2. If the `/etc/nvidia/gridd.conf` file does not already exist, create it by copying the supplied template file `/etc/nvidia/gridd.conf.template`.
3. As root, edit the `/etc/nvidia/gridd.conf` file to set the `EnableUI` option to `TRUE`.
4. Start the `nvidia-gridd` service.

```
# sudo service nvidia-gridd start
```

When **NVIDIA X Server Settings** is restarted, the **Manage License** option is now available.

Status

Open

3.21. Licenses remain checked out when VMs are forcibly powered off

Description

NVIDIA vGPU software licenses remain checked out on the license server when non-persistent VMs are forcibly powered off.

The NVIDIA service running in a VM returns checked out licenses when the VM is shut down. In environments where non-persistent licensed VMs are not cleanly shut down, licenses on the license server can become exhausted. For example, this issue can occur in automated test environments where VMs are frequently changing and are not guaranteed to be cleanly shut down. The licenses from such VMs remain checked out against their MAC address for seven days before they time out and become available to other VMs.

Resolution

If VMs are routinely being powered off without clean shutdown in your environment, you can avoid this issue by shortening the license borrow period. To shorten the license borrow period, set the `LicenseInterval` configuration setting in your VM image. For details, refer to *Virtual GPU Client Licensing User Guide*.

Status

Closed

Ref.

1694975

Notice

This document is provided for information purposes only and shall not be regarded as a warranty of a certain functionality, condition, or quality of a product. NVIDIA Corporation ("NVIDIA") makes no representations or warranties, expressed or implied, as to the accuracy or completeness of the information contained in this document and assumes no responsibility for any errors contained herein. NVIDIA shall have no liability for the consequences or use of such information or for any infringement of patents or other rights of third parties that may result from its use. This document is not a commitment to develop, release, or deliver any Material (defined below), code, or functionality.

NVIDIA reserves the right to make corrections, modifications, enhancements, improvements, and any other changes to this document, at any time without notice.

Customer should obtain the latest relevant information before placing orders and should verify that such information is current and complete.

NVIDIA products are sold subject to the NVIDIA standard terms and conditions of sale supplied at the time of order acknowledgement, unless otherwise agreed in an individual sales agreement signed by authorized representatives of NVIDIA and customer ("Terms of Sale"). NVIDIA hereby expressly objects to applying any customer general terms and conditions with regards to the purchase of the NVIDIA product referenced in this document. No contractual obligations are formed either directly or indirectly by this document.

NVIDIA products are not designed, authorized, or warranted to be suitable for use in medical, military, aircraft, space, or life support equipment, nor in applications where failure or malfunction of the NVIDIA product can reasonably be expected to result in personal injury, death, or property or environmental damage. NVIDIA accepts no liability for inclusion and/or use of NVIDIA products in such equipment or applications and therefore such inclusion and/or use is at customer's own risk.

NVIDIA makes no representation or warranty that products based on this document will be suitable for any specified use. Testing of all parameters of each product is not necessarily performed by NVIDIA. It is customer's sole responsibility to evaluate and determine the applicability of any information contained in this document, ensure the product is suitable and fit for the application planned by customer, and perform the necessary testing for the application in order to avoid a default of the application or the product. Weaknesses in customer's product designs may affect the quality and reliability of the NVIDIA product and may result in additional or different conditions and/or requirements beyond those contained in this document. NVIDIA accepts no liability related to any default, damage, costs, or problem which may be based on or attributable to: (i) the use of the NVIDIA product in any manner that is contrary to this document or (ii) customer product designs.

No license, either expressed or implied, is granted under any NVIDIA patent right, copyright, or other NVIDIA intellectual property right under this document. Information published by NVIDIA regarding third-party products or services does not constitute a license from NVIDIA to use such products or services or a warranty or endorsement thereof. Use of such information may require a license from a third party under the patents or other intellectual property rights of the third party, or a license from NVIDIA under the patents or other intellectual property rights of NVIDIA.

Reproduction of information in this document is permissible only if approved in advance by NVIDIA in writing, reproduced without alteration and in full compliance with all applicable export laws and regulations, and accompanied by all associated conditions, limitations, and notices.

THIS DOCUMENT AND ALL NVIDIA DESIGN SPECIFICATIONS, REFERENCE BOARDS, FILES, DRAWINGS, DIAGNOSTICS, LISTS, AND OTHER DOCUMENTS (TOGETHER AND SEPARATELY, "MATERIALS") ARE BEING PROVIDED "AS IS." NVIDIA MAKES NO WARRANTIES, EXPRESSED, IMPLIED, STATUTORY, OR OTHERWISE WITH RESPECT TO THE MATERIALS, AND EXPRESSLY DISCLAIMS ALL IMPLIED WARRANTIES OF NONINFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE. TO THE EXTENT NOT PROHIBITED BY LAW, IN NO EVENT WILL NVIDIA BE LIABLE FOR ANY DAMAGES, INCLUDING WITHOUT LIMITATION ANY DIRECT, INDIRECT, SPECIAL, INCIDENTAL, PUNITIVE, OR CONSEQUENTIAL DAMAGES, HOWEVER CAUSED AND REGARDLESS OF THE THEORY OF LIABILITY, ARISING OUT OF ANY USE OF THIS DOCUMENT, EVEN IF NVIDIA HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Notwithstanding any damages that customer might incur for any reason whatsoever, NVIDIA's aggregate and cumulative liability towards customer for the products described herein shall be limited in accordance with the Terms of Sale for the product.

VESA DisplayPort

DisplayPort and DisplayPort Compliance Logo, DisplayPort Compliance Logo for Dual-mode Sources, and DisplayPort Compliance Logo for Active Cables are trademarks owned by the Video Electronics Standards Association in the United States and other countries.

HDMI

HDMI, the HDMI logo, and High-Definition Multimedia Interface are trademarks or registered trademarks of HDMI Licensing LLC.

OpenCL

OpenCL is a trademark of Apple Inc. used under license to the Khronos Group Inc.

Trademarks

NVIDIA, the NVIDIA logo, NVIDIA GRID, NVIDIA GRID vGPU, NVIDIA Maxwell, NVIDIA Pascal, NVIDIA Turing, NVIDIA Volta, GPUDirect, Quadro, and Tesla are trademarks or registered trademarks of NVIDIA Corporation in the U.S. and other countries. Other company and product names may be trademarks of the respective companies with which they are associated.

Copyright

NVIDIA Corporation | 2788 San Tomas Expressway, Santa Clara, CA 95051
http://www.nvidia.com